

demographic characteristics.¹⁵⁹ As a result, it can be said that fMRI studies within the first two decades of the twenty-first century have placed the hysteria patient into a decidedly somatic framework. Not only have these studies aimed to determine a neurophysiological basis of hysteria, but they have also judged the patients' adequacy as a potential study participant by focusing exclusively on their quantifiable physical symptoms. Notably, both aspects of this purely somatic framing of the present-day hysteria patient as an experimental subject are curiously reminiscent of how Charcot had approached his patients more than a century earlier.¹⁶⁰

To summarise, my analysis has shown that, on the whole, the inclusion in an fMRI study has tended to strip hysteria patients of the messy multisymptomatic materiality of their disease while also detaching them from individual life events that might have given rise to their symptoms. To speak with Latour,¹⁶¹ the transformation of hysteria patients into experimental subjects in the first two decades of the twenty-first century has entailed the amplification of those aspects of their disorder that were judged to have a shared neural basis and could thus be addressed adequately by the fMRI measurement. At the same time, the patient selection has also involved the reduction of the idiosyncratic features that might have had the potential to skew the results by introducing unwanted variability into the imaging data. I thus argue that from the perspective of an fMRI experimental setup, hysteria patients are viewed as contingent variables. In other words, hysteria patients are treated as products of chance that need to be disciplined through sampling to meet the technological requirements of fMRI. Through such disciplining that underpins the inclusion into an fMRI study, each hysteria patient becomes part of the chain of transformations on whose consistency the meaning of the resulting functional brain maps hinges. Hence, we have seen that the participant selection, together with the choice of the experimental task and the conditions of its implementation, play crucial roles in making the hysterical symptom measurable through fMRI.

3.2 Measurement: Translating the Active Brain into Imaging Data

Having recruited the experimental subjects and programmed the task implementation, researchers can finally start to collect imaging data by scanning each subject's brain separately. For this purpose, the subject enters the scanner room and lays face upwards on the machine's moveable table.¹⁶² Here, she receives earplugs and headphones to

159 See, e.g., Espay et al., "Functional Tremor," 183; and Hassa et al., "Motor Control," 144.

160 For a discussion of Charcot's somatic framing of his patients' emotional states and memories of traumatic experiences, see, in particular, sections 1.1.3, 1.2.2, and 1.3.2.

161 Latour, *Pandora's Hope*, 70–71.

162 The following description is based on my experience of participating as a healthy control subject in 2012 in two fMRI studies conducted at the Charité Campus Mitte Berlin. Moreover, on multiple occasions in 2014 and 2015, I sat with researchers in a control room of the fMRI scanning facility at the Department of Psychiatry and Psychotherapy, Charité Campus Mitte Berlin, while they were

protect her from high levels of acoustic noise that characterise the measurement.¹⁶³ After that, she is handed a button box or a joystick with which she will respond to the task. The subject's head is then placed into a cage-like plastic cylinder called a head coil and firmly fixed within it with paddings. The subject is instructed to remain as still as possible during the entire measurement. Finally, the table is moved into the measurement position in the middle of the scanner's bore.

During the measurement, researchers sit in the adjacent console room, which is connected to the scanner room via a large observation window. In the pauses between different stages of the measurement, researchers can communicate with the subject via an intercom. They operate the computers that simultaneously control the scanning procedure and the concurrent exposure of the subject to pre-programmed experimental stimuli. The measurement begins with the acquisition of a low-resolution structural scan called the localiser, which appears on the computer screen within a minute (fig. 3.1). This image shows a vertical section of the subject's brain within the skull. By providing the information about the position of the subject's head within the scanner, this image allows researchers to optimise the location of subsequent scans.¹⁶⁴ In the next eight to ten minutes, researchers collect a high-resolution structural scan (fig. 3.2). The structural scan provides information about the subject's brain anatomy and, as we will see later, plays an important role in the analysis of functional scans.

Figure 3.1. View of a computer screen showing a localiser scan.




collecting imaging data from healthy control subjects. I am grateful to Torsten Wüstenberg for making this possible.

163 For details on the acoustic noise, see Huettel, Song, and McCarthy, *Imaging*, 54.

164 Ashby, *Statistical Analysis*, 4.

Only after structural scanning has been completed does the experiment in itself begin. Over the next twenty to thirty minutes, the subject carries out pre-programmed task instructions while, in a synchronised and fully automated process, the scanner generates functional imaging data.¹⁶⁵ During the measurement, researchers can view the incoming fMRI data, which by this point have already undergone several algorithmic transformations.¹⁶⁶ To enable their viewing, fMRI data are automatically visualised on the computer screen as fuzzy grey-scale images of brain slices (fig. 3.3). Crucially, however, for reasons I will discuss in the following sections, by submitting these images to visual inspections, researchers are unable to determine whether and in which anatomical locations their experimental manipulation induced brain activity of interest. Hence, to make judgments about the task-induced brain activity, researchers must first transform the essentially illegible fMRI data into functional brain maps (see figs. 3.12–3.15) through statistical analysis. Since the conditions of data acquisition have epistemic implications for functional brain maps that are constructed from them—and, in turn, inform the kind of knowledge that an fMRI study can produce about hysteria—we must examine how imaging data are generated.


Figure 3.2. Spatial visualisation of an experimental subject's structural imaging data.



165 Functional (imaging) data, fMRI data, functional/fMRI scans, functional/fMRI images and functional MR images are synonymous terms that I interchangeably use throughout the chapter.

166 This will become evident in the course of my analysis in the following sections.

Figure 3.3. Spatial visualisation of BOLD fMRI data obtained through the EPI (echo-planar imaging) sequence.



In the following three sections, I will explore how the measurement creates a referential link between the imaging data and the living brain. We will see that an MR scanner cannot generate functional imaging data in a single step. Instead, by inducing the eponymous magnetic nuclear resonance, the scanner manipulates the subatomic particles in the brain tissue into producing a physical quantity called a magnetic resonance (MR) signal.¹⁶⁷ This signal forms the basis for both structural and functional imaging data. But what exactly happens under the hood during the scanning? Through which operations does the measurement bridge the successive gaps between the active brain, the MR signal, and the different types of imaging data that appear on the screen? To what extent can researchers shape this computer-controlled process?

My step-by-step analysis will answer these questions by tracing, first, how MR signals are generated; second, how these signals are transformed into spatially configured images of brain slices; and finally, how a vast amount of fMRI slices are collected at different time points. I will particularly focus on elucidating the theoretical assumptions that are implicitly built into the imaging data as a precondition for making the active human brain measurable and visualisable. Moreover, by the end of my analysis, it will become apparent that while fMRI enables present-day hysteria researchers to obtain physiologically far more proximate access to their patients' brain activity than Charcot ever had, this imaging technology also imposes on researchers an entirely different way of both working on and with images.

¹⁶⁷ Huettel, Song, and McCarthy, *Imaging*, 25, 38.

3.2.1 Generating the Initial Inscriptions

The majority of fMRI studies in general and all fMRI studies of hysteria discussed in this book have employed the BOLD fMRI method.¹⁶⁸ As mentioned in chapter 2, BOLD stands for the blood-oxygen-level-dependent contrast. This method relies on the well-established experimental fact that neural activity is accompanied by a set of mutually interrelated metabolic and physiological changes in the brain.¹⁶⁹ In short, active neurons have heightened energy requirements that are fulfilled through the oxidation of glucose. To enable the oxidation, glucose and oxygen must be supplied to the active brain regions through a local increase in blood flow. For this purpose, oxygen is attached to haemoglobin, a protein contained in the blood. Having delivered oxygen, haemoglobin binds carbon dioxide, the waste product of oxidation, and transports it out of the brain. Haemoglobin carrying oxygen is not magnetic, whereas deoxygenated haemoglobin has strong magnetic properties.¹⁷⁰ Thus, blood has measurably different magnetic properties depending on the relative concentration of deoxygenated haemoglobin that it contains. Crucially, such changes in the relative concentration of deoxygenated haemoglobin serve as a proxy for the correlated neural activity and provide the basis for the BOLD fMRI method.¹⁷¹

However, the MR scanner cannot directly measure the changing concentrations of deoxygenated haemoglobin. Hence, magnetic resonance imaging deploys an indirect approach and utilises instead the fact that hydrogen nuclei are chemically bound in water and thus abundant in the brain.¹⁷² Each nucleus consists of a single proton, a positively charged subatomic particle, whose behaviour is governed by the laws of quantum mechanics.¹⁷³ The behaviour of each proton is characterised by a quantum mechanical property called the nuclear spin that, in an analogy to classical mechanics, can approximately be imagined as rotation around its own axis.¹⁷⁴ Since they are electrically charged, as a consequence of this 'self-rotation,' protons behave as magnetic dipoles and produce minuscule magnetic fields.

Relying on the quantum mechanics laws, the scanner manipulates the nuclear magnetic behaviour of hydrogen nuclei and then measures the thus induced effects in the form of MR signals. Depending on the exact parameters of the manipulation,

168 Huettel, Song, and McCarthy, 131. See also Ogawa et al., "Oxygenation-Sensitive."

169 Logothetis, "Neural Basis," 1004.

170 Logothetis, 1009.

171 Seiji Ogawa discovered the BOLD method in 1990 by experimentally demonstrating that an MR scanner can be used to generate images, which register the variations in the concentration of deoxygenated haemoglobin. See Ogawa et al., "Oxygenation-Sensitive," 68–78. Yet, Ogawa's experiments were performed on rats. For the first application of the BOLD method to the imaging of the living human brain, see Kwong et al., "Human Brain Activity."

172 Haacke et al., *Imaging*, 3.

173 Haacke et al., 3.

174 To be exact, nuclei do not actually rotate around their axis. Instead, nuclear spin is a theoretical construct developed in 1924 by the Austrian physicist Wolfgang Pauli to explain certain experimental findings. Haacke et al., 11. As an intrinsic property of subatomic particles, spin can adequately be described only by mathematical formulations of quantum mechanics. See *ibid.*, 71–73.

the resulting signals provide information about differences in the nuclear magnetic behaviour of various hydrogen nuclei (i.e., protons) across the brain.¹⁷⁵ As my analysis in the following will show, on the one hand, such signals underlie the construction of structural images that entail information about brain anatomy. On the other hand, such signals also enable the production of BOLD fMRI images that contain information about the relative concentrations of deoxygenated haemoglobin across the brain. Researchers, in turn, use the information obtained through such images to make judgments about both static and dynamic features of their experimental subjects' brains.

Yet, how does the black-boxed measurement construct the referential relationship between the different types of imaging data, on the one hand, and the brain's static as well as dynamic features, on the other? To answer this question, we must examine the process through which the scanner first generates and then samples MR signals. The details of this process are complex and best described by mathematical equations.¹⁷⁶ In this section, my simplified account will selectively highlight those aspects of this complex process that are relevant to my subsequent discussion of the fMRI-based chains of references in contemporary hysteria research.¹⁷⁷

The measurement starts by placing the head of the experimental subject into the MR scanner's static magnetic field, whose strength is typically 1.5 or 3 Tesla (T).¹⁷⁸ Under normal conditions, the magnetic fields of individual hydrogen nuclei are distributed randomly in the brain, hence cancelling each other out so that the human brain as a whole has no significant magnetic properties.¹⁷⁹ However, the scanner's magnetic field interacts with the magnetic fields of individual protons (i.e., hydrogen nuclei). Due to this interaction, the magnetic fields of individual protons align in the direction of the scanner's field, thus causing the brain to 'generate' its own magnetic field.¹⁸⁰ To make the magnetic properties of the thus aligned protons measurable, in the next step,

175 Huettel, Song, and McCarthy, *Imaging*, 121. For the sake of simplicity, in the rest of the text, I will refer to hydrogen nuclei as protons.

176 For the precise mathematical description of all stages of the measurement, see Haacke et al., *Imaging*, 1–380.

177 I am using the term chain of references in Latour's sense. See Latour, "More Manipulation," 348.

178 The field of a 3T scanner is approximately 60,000 times stronger than the Earth's magnetic field. See Huettel, Song, and McCarthy, *Imaging*, 3. Stronger scanners are also available but have so far been rarely used in hysteria research. Two exceptions are Espay et al., "Functional Dystonia"; and Espay et al., "Functional Tremor." Both studies used a 4T scanner.

179 For a detailed description, see Huettel, Song, and McCarthy, *Imaging*, 59–60.

180 Simply put, when exposed to the scanner's magnetic field, the previously randomly distributed protons are forced to re-orientate. As a result of this re-orientation, they start a gyroscopic motion around the axis determined by the direction of the external field. This rotational movement is called precession. The frequency of this motion depends on the strength of the external field. Moreover, according to the quantum mechanics laws, protons are allowed to occupy only two distinct orientations in an external magnetic field—either a low-energy state parallel to the external field or a high-energy antiparallel state. Due to a higher probability of protons occupying the low-energy state, a larger number of them align themselves parallel to the scanner's field. After the protons have reached this equilibrium alignment, their individual magnetic fields add up to generate the brain's magnetic field. The resulting field is parallel to that of the scanner and is designated by the physical value called the net magnetisation. See Huettel, Song, and McCarthy, 59–64.

the brain is exposed to an electromagnetic wave that oscillates in the radio-frequency (RF) range. Called RF pulse, this wave is a dynamic magnetic field that adds energy to the aligned protons.¹⁸¹ In the process referred to as excitation, some protons absorb this energy and change their orientation in the scanner's field.¹⁸² As a result of this intervention, the brain's magnetic field is temporarily brought out of balance.

After the RF pulse has been switched off, the protons gradually return to the initial orientation during the process called relaxation.¹⁸³ Depending on the local nuclear magnetic properties of the brain tissue, various protons undergo the process of relaxation at different speeds.¹⁸⁴ While returning to its initial state, each proton re-emits the energy it had absorbed during the excitation. It thereby generates an electromagnetic wave of the same frequency as the RF pulse.¹⁸⁵ The waves from different protons add up into a cumulative MR signal. The intensity of an MR signal—which depends on the density of the protons that relax at the same speed—is registered in a numerical form by the detectors located in the volume coil around the subject's head.¹⁸⁶ In specialist terms, the registered signal intensities are called raw data and serve as the basis for generating all types of imaging data.¹⁸⁷ But since the process of their transformation into images is automated, raw data are not directly accessible in contemporary scanners.

It follows from my description that the cumulative signal intensities registered by the detectors—i.e., the raw data—are materialised traces of a physical interaction between the hydrogen nuclei in the subject's brain and a specific combination of the scanner's static and dynamic magnetic fields. Consequently, the resulting signals have an indexical relation to the brain. As the initial inscriptions generated by the measurement, these signals provide the material basis for an entire cascade of transformations, which, by the end of the experiment, result in the production of functional brain maps. Throughout this chapter, I will trace how each transformation in this cascade aims to amplify the information of interest in the fMRI data while

181 An electromagnetic wave carries an amount of energy determined by its frequency. Huettel, Song, and McCarthy, 64.

182 According to quantum mechanical rules, the lined-up protons can only absorb a certain amount of energy called a quantum, which is equal to the energy difference between their parallel and antiparallel states in the external magnetic field. Such quanta of energy correspond to specific, so-called resonant frequencies of the electromagnetic pulse and are equal to the precession frequency of the proton at a given strength of the scanner's magnetic field. When exposed to the RF pulse, multiple nuclei absorb the energy corresponding to their resonant frequencies and jump from the parallel to the antiparallel orientation. This leads to a reversible change in the brain's net magnetisation. For details, see Huettel, Song, and McCarthy, 64.

183 During this period, the brain's magnetisation, induced through the scanner's external field, gradually regains its initial value as well as its orientation along the main axis of the scanner's magnetic field. Huettel, Song, and McCarthy, 85–87.

184 Huettel, Song, and McCarthy, 66–67. As we will see shortly, this plays a principal role in the imaging.

185 This physical process during which atomic nuclei that are located in an external magnetic field first absorb the energy of an electromagnetic pulse and then re-emit it is the eponymous nuclear magnetic resonance. Huettel, Song, and McCarthy, 18.

186 Haacke et al., *Imaging*, 5–6.

187 I will discuss how MR signals are transformed into images in the subsequent section.

preserving an unbroken link to the underlying indexical MR signals. Thus, we will see that even if no longer directly accessible, the indexical raw data are the very foundation of the functional maps' ability to refer to and thus produce scientific insights into the workings of the active human brains.

So far, we have examined how MR signals are created. This leaves us with the question of how these signals are brought into referential relations to brain anatomy and activity, respectively. For this purpose, the measurement deploys the fact that during relaxation—the period over which the excited protons return to their initial state—the intensity of the MR signal changes under the influence of multiple factors that characterise the protons' local environments.¹⁸⁸ In mathematical terms, this change in the signal intensity is described by a quantum mechanical equation, in which three different time constants designate the factors of critical importance for the imaging.¹⁸⁹ Whereas the time constant T_1 specifies the rate with which protons return to the state that preceded excitation, T_2 describes the signal decay due to the mutual interactions among the protons.¹⁹⁰ These two constants are determined by the inherent physiological properties of the tissue at a given strength of the scanner's magnetic field.¹⁹¹ Since they have distinct values for different types of brain tissue, these two constants play key roles in structural imaging. By contrast, T_2^* , the time constant crucial for functional imaging, describes the rate of the signal decay due to the combined effects of intrinsic tissue properties and the presence of irregularities in the local magnetic fields.¹⁹²

Crucially, the quantum mechanical signal equation provides the theoretical framework that informs the entire imaging process. By relying on this equation, a specific time point during the relaxation can be determined that allows the measurement to highlight the effect of the chosen time constant on the signal's intensity while simultaneously minimising the effects of the other nuclear magnetic properties. This is achieved by varying the time intervals between the RF pulse excitation and the signal collection (i.e., the echo time), as well as the intervals between two successive RF pulses (i.e., the repetition time).¹⁹³ To put it more plainly, through targeted variations in the timing of the data sampling, the scanner can selectively highlight the inscription of the protons' chosen magnetic property into the signal, thus allowing the construction of various types of imaging data. My analysis thus makes evident that signal sampling has a crucial semantic role—it is not only the way the data are generated but also how they are collected that determines their informational content.

In contemporary scanners, the details of the complex sampling processes deployed to generate different types of imaging data are called pulse sequences and are black-

188 The signal intensity depends "on the (proton) spin density, the so-called T_1 and T_2 relaxation times, and on other physical parameters of the tissue such as diffusion, perfusion or velocity (e.g. blood flow)." Logothetis, "Neural Basis," 1006.

189 This is the famous Bloch equation. See Haacke et al., *Imaging*, 8–9.

190 Huettel, Song, and McCarthy, *Imaging*, 66.

191 Huettel, Song, and McCarthy, 66–67.

192 Huettel, Song, and McCarthy, 131.

193 For more details, see Huettel, Song, and McCarthy, 122–52.

boxed behind the software.¹⁹⁴ Researchers can choose among various pulse sequences that issue automated commands to the scanner with which relative timing to execute the operations underlying the signal generation and sampling.¹⁹⁵ Each such selective intervention measures a distinct physical property and produces a particular type of MR image that spatially encodes the information about how the relative values of the physical property chosen change across the brain.¹⁹⁶ In specialist terms, such different types of images are referred to as image contrasts.¹⁹⁷ Two types of contrasts are of interest to us.

If the signal is acquired to highlight the effect of the T_1 constant, the results are T_1 -weighted images, which are the most commonly used structural data in fMRI studies (see fig. 3.2).¹⁹⁸ These images encode the spatial distribution of the relative T_1 values across the brain. Upon finished acquisition, they are automatically visualised as grey-scale brain slices in which the brightest parts refer to white matter, intermediate to grey matter, whereas the fluid-filled cavities are shown in black. In fact, due to the way in which the tissue types are accorded relative grey values, at a superficial glance, T_1 images visually resemble black-and-white photographs of a dissected brain. Yet, my analysis above has underscored the distinctly non-mimetic character of these images. We have seen that the referential link of such images to the brain's anatomy rests on the empirically established fact that there is a one-to-one relationship between the T_1 values and the intrinsic physiological properties of different brain tissues.¹⁹⁹

Conversely, the most widely used functional imaging data are referred to in specialist terms as T_2^* -weighted images (see fig. 3.3).²⁰⁰ Their production relies on the same physical principles and mathematical models as T_1 images. However, T_2^* images encode a different type of information. T_2^* contrast is produced by the pulse sequence that highlights the effects of the local magnetic field irregularities on the loss of the MR signal intensity.²⁰¹ At first glance, it may seem counterintuitive that signal disturbances can provide useful information about the brain. However, this type of targeted measurement makes clever use of a well-established experimental fact. As mentioned at the beginning of this section, heightened neural activity leads to changes

194 See Huettel, Song, and McCarthy, 3, 92–93.

195 Huettel, Song, and McCarthy, 122–52.

196 “In MRI the signals are arbitrarily scaled and there are no units.” Jenkinson and Chappell, *Neuroimaging Analysis*, 25. Hence, the information of interest in the images is not expressed in absolute values. Instead, it is conveyed through the relative differences in the signal intensities across the image.

197 An MR scanner can be used to generate a wide variety of both static and motion contrasts. Static contrasts are “sensitive to the type, number, relaxation properties, and local environment” of atomic nuclei. Motion contrasts are sensitive to the protons’ movement. The different contrasts provide information about the brain's anatomy, neural activity, blood flow, water diffusion, and other aspects of interest. Jenkinson and Chappell, 121.

198 Jenkinson and Chappell, 128.

199 Jenkinson and Chappell, 126–28. In a T_1 -contrast image, a tissue with a short T_1 value—such as white matter—appears bright, whereas a tissue with a very long T_1 value—such as cerebrospinal fluid—appears black.

200 Jenkinson and Chappell, 131.

201 Jenkinson and Chappell, 131.

in the relative concentration of deoxygenated haemoglobin in the capillaries in the vicinity of the active brain areas.²⁰² Since varying concentrations of deoxygenated haemoglobin in blood have different magnetic properties, they produce dynamic local irregularities in the magnetic field within the brain.²⁰³ These irregularities have a measurable effect on the decay rate of the MR signal intensity and are, therefore, used as the basis for generating T_2^* images.

Upon acquisition, T_2^* images are also automatically rendered as grey-scale brain slices that can be inspected on a computer screen. Significantly, the different hues of grey in these images encode the relative differences in the MR signal decay rates due to the local magnetic field disturbances, which, in turn, are caused by various levels of deoxygenated haemoglobin across the brain.²⁰⁴ In other words, T_2^* images encode the blood-oxygen-level-dependent (BOLD) contrast and thus enable researchers to make inferences about regionally specific brain activity. Yet, as underscored by my analysis, the production of these images is anything but straightforward. Instead, it requires bridging multiple gaps to reach the phenomenon of interest to which the images refer. As we have seen, the relative differences in the signal intensity encoded in the images are used as a proxy for various levels of oxygen concentration in the blood. These concentrations are “determined by the balance of supply (blood flow) and demand (extraction by tissue) of oxygen,” which, in turn, serve as indicators of the correlated neural activity.²⁰⁵

To sum up, in this section, I have shown that the generation of indexical MR signals is only the first step in constructing the referential link between the active brain and different types of imaging data. This initial step is necessarily followed by a specifically tailored sampling procedure that selectively inscribes into the MR signal those nuclear magnetic properties, which serve as indicators of the brain's pertinent static and dynamic features at the macro level. Consequently, the informational content of the resulting images is entirely predicated on the concrete conditions of this active intervention. These conditions include the strength and quality of the scanner's magnetic field, various parameters of the RF excitation, and the details of the pulse sequences chosen. Hence, although the operations underlying the production of MR images are automated and thus black-boxed, their traceability remains a key precondition for the epistemic validity of the fMRI imaging data in the scientific context. For this reason, researchers are required to report in detail in their published papers not only which scanner they used—by stating its field strength, the manufacturer and the model—but also the exact pulse sequences with which they generated each type of imaging data in their study.

202 Logothetis, “Neural Basis,” 1004.

203 For details, see Huettel, Song, and McCarthy, *Imaging*, 193–96, 198–200.

204 See Ogawa et al., “Oxygenation-Sensitive,” 68–78.

205 Ogawa et al., “Blood Oxygenation,” 9872.

3.2.2 Constructing the Spatiality of the Imaging Data

So far, we have examined how the brain's nuclear magnetic properties of interest are inscribed into the imaging data. We have also discussed how, upon finished acquisition, the resulting data are automatically visualised as grey-scale brain slices that display a spatial distribution of the properties measured. The spatial information contained in the data is of fundamental importance because what makes fMRI an imaging technology is its ability to localise the neural activity of interest within the subject's brain. Yet, how does the measurement construct the spatiality of the imaging data? Moreover, how does the image space relate to the physical space of the subject's brain? To answer these questions, in this section, I will analyse the black-boxed operations that underlie the creation of functional and structural images from MR signals.

The first step in constructing the imaging data's spatiality entails generating a cumulative MR signal in a way that enables the subsequent reconstruction of the relative spatial locations from which the individual contributions making up that signal had originated.²⁰⁶ Called spatial encoding, this process hinges on the introduction of controlled changes into the uniform magnetic field of the scanner during the measurement. This is achieved by employing additional magnetic fields called gradients, whose strength changes linearly in one direction in a known way.²⁰⁷ Contemporary MR scanners implement a standard configuration in which the main magnetic field is superimposed by three mutually orthogonal time-dependent linear gradients oriented along the axes of the 3D Cartesian coordinate system.²⁰⁸ By convention, the z-axis has a foot-to-head direction, whereas the x-y plane is perpendicular to the z-axis.²⁰⁹

It should be noted that the Cartesian coordinate system is never visualised in a perceptible way during the measurement. Instead, it is used as an abstract framework whose function is to mathematically describe the spatial configuration of the time-dependent gradients that intervene in the physical space of the brain. In a series of steps that follow an RF pulse excitation, a particular combination of the gradients along the axes of the 3D Cartesian coordinate system is sequentially switched on and off.²¹⁰

206 Lauterbur, "Appendix A," 235.

207 Lauterbur, 236. Interestingly, the phenomenon of nuclear magnetic resonance that underlies the entire MRI imaging was discovered and described between the late 1930s and mid-1940s. See, e.g., Huettel, Song, and McCarthy, *Imaging*, 15–18. However, it was only in the early 1970s that the American chemist Paul Lauterbur introduced the idea of spatial encoding by means of magnetic field gradients, thus enabling the translation of MR signals into images that visualise internal spatial structures of opaque objects. Lauterbur thus paved the way for the development of MRI as an imaging technology, for which he received the Nobel Prize in Physiology or Medicine in 2003. For Lauterbur's seminal article published in *Nature* in 1973, see Lauterbur, "Image Formation." For his Nobel Prize acceptance speech, see Lauterbur, "All Science is Interdisciplinary."

208 Huettel, Song, and McCarthy, *Imaging*, 90–91.

209 Huettel, Song, and McCarthy, 90–91. The main field of the scanner is oriented along the z-axis.

210 The rapid, sequential switching on and off of the magnetic gradients makes the measurement loud. For details, see Huettel, Song, and McCarthy, 91–96. My analysis in this section focuses on the so-called Cartesian acquisition method in which individual data points are sampled along the axes of the Cartesian coordinate system. The Cartesian acquisition is "widely accepted as the

Through this intervention, the brain is mathematically segmented into a virtual 3D grid that consists of many individual cubes. As a result, the physical space of the brain is coordinatised—the location of each virtual cube is designated by a unique combination of spatial coordinates. Within these virtual cubes, the gradients physically interact with the local protons, forcing them to behave differently depending on their relative spatial locations within the scanner's coordinate system.²¹¹ Through this targeted intervention, the gradients manipulate the protons within each virtual cube into producing a specific MR signal into which two different kinds of information are encoded. In addition to the information about the local brain physiology, the resulting signal also contains information about the protons' relative spatial locations within the brain.

In the 3D data subsequently reconstructed from the signals, virtual cubes are represented by voxels (i.e., 'volume elements'), which are the 3D equivalent of pixels (i.e., 'picture elements'). Each voxel is designated by a particular set of the Cartesian coordinates (x, y, z) that link it to the physical location in the brain that had been labelled by the same set of coordinates during the acquisition.²¹² In effect, through the intervention of the gradients, the Cartesian coordinates are first encoded into the signal and then, during the image reconstruction, subsequently built into the imaging data. The original coordinates that are thus attributed to each voxel of the reconstructed 3D image jointly comprise the so-called "native space" of the image.²¹³ Importantly, the resulting voxels are the smallest visual elements of a 3D image. Therefore, their size determines the spatial resolution of the imaging data. That size can vary from less than a cubic millimetre for structural to several cubic millimetres for functional images.²¹⁴ Depending on the manner of visualisation chosen by researchers or, in some cases, hard-coded into the software, each voxel can be assigned a single numerical value, colour, or shade of grey.

This succinct description demonstrates that the Cartesian coordinate system plays a pivotal role in bridging the gaps between the subject's brain, the MR signal and the

standard techniques." Block et al., "Clinical Use," 87. Recently, non-Cartesian methods such as spiral and radial imaging have been developed. In these methods, the gradients are modulated in non-linear ways during the measurement. See Block et al., 88–89; and Huettel, Song, and McCarthy, *Imaging*, 148–50. But even in the non-Cartesian methods, the raw signal data thus sampled have to be "interpolated back onto a Cartesian grid" before the automated algorithms "can be used to reconstruct the image." Huettel, Song, and McCarthy, 149. See also Block et al., "Clinical Use," 90. Hence, even in the non-Cartesian methods, the Cartesian coordinate system plays an important role in the construction of the imaging data's spatiality.

211 Due to the use of the gradients, the strength of the effective magnetic field differs across the virtual cubes, depending on their respective locations within the scanner's coordinate system. As mentioned previously, the resonant frequencies of lined-up protons are determined by the cumulative strength of the imposed magnetic fields. Hence, through the introduction of the gradients, the location of each cube is labelled with a distinct resonant frequency. The different resonant frequencies can be mathematically reconstructed from the cumulative MR signal, thus linking a distinct frequency to a particular location within the brain. See Lauterbur, "Appendix A," 235.

212 Poldrack, Mumford, and Nichols, *Handbook*, 17.

213 Poldrack, Mumford, and Nichols, 17.

214 See, e.g., Morris et al., "Avoidance Learning," 288–89.

images by making the spatial distribution of the brain anatomy and activity visualisable. Basically, the measurement maps the mathematically defined spatial relations (i.e., the Cartesian coordinates) onto the brain's opaque space and encodes them into the signals sampled. Since the Cartesian coordinate system allows a seamless transformation from a mathematical to a visual representation of space, the spatial structure that was numerically encoded into the signals during the measurement can subsequently be translated into the geometric space of an MRI image.²¹⁵ The brain's continuous physical space is thus mathematised and recast into a discretised 3D space of functional and structural images.²¹⁶ It can, therefore, be argued that the deployment of the Cartesian coordinate system enables researchers to visually configure the otherwise indiscernible space of the brain enclosed within the skull.

Moreover, because the Cartesian coordinate system spatially structures the physical interaction between the gradients and the brain, it establishes a referential link between the voxels in the image and the physical locations within the subject's brain. Thus, the indexicality of the imaging data does not hinge only on the nature of the signal, as discussed previously, but also on the use of the Cartesian coordinate system. Due to the use of the coordinate system, the signal is generated to provide indexical information about the relative spatial distribution of the brain's nuclear magnetic property of interest. In fact, as we will see later in the chapter, the Cartesian coordinate system remains the central organising principle of the imaging data's spatial features and the conveyor of their referential quality through all stages of data analysis.

Significantly, the construction of the 3D image space is a complex process that requires an entire chain of well-nested operations. Since three-dimensional imaging is slow and computationally challenging, most structural and all functional imaging is performed by acquiring a succession of 2D cross-sectional slices one at a time.²¹⁷ A collection of slices that add up to cover the entire space of the brain is jointly called a volume. Notably, a volume "is considered to be a single image" in the neuroimaging context.²¹⁸ However, it is worth mentioning that such a 3D 'single image' is not visually graspable in its entirety for its human users, nor can it be viewed at a single glance.²¹⁹ Yet, the point I want to make in the following is that even the acquisition of a single 2D slice, which is only a fragment of a 3D image, is not a one-step procedure but instead necessitates multiple transformations.

-
- 215 As the philosopher of science Bas van Fraassen has pertinently pointed out, by applying algebraic equations to geometric figures, René Descartes created his coordinate-based method in which every location in space could be represented by a set of numerical values. See van Fraassen, *Scientific Representation*, 66–67. The use of the Cartesian coordinate system thus makes possible not only the operation of structuring an arbitrary space by means of numerical values and functions (as in the process of MRI acquisition) but also the opposite process of translating abstract numerical into corresponding spatial relations (as during the reconstruction of images from MR signals).
- 216 My analysis is in line with Michael Lynch's discussion on the role of mathematisation in the production of scientific visibility. See Lynch, "Material Form of Images," 37–66.
- 217 Huettel, Song, and McCarthy, *Imaging*, 91.
- 218 Jenkinson and Chappell, *Neuroimaging Analysis*, 23.
- 219 This will become more apparent in section 3.5.1 when I turn to analysing how researchers visually inspect 3D imaging data.

First, to select a slice in the x-y plane, a gradient along the z-axis is applied simultaneously with an RF pulse.²²⁰ This intervention makes only a small proportion of the protons sensitive to the RF pulse, thus resulting in the selective excitation of a particular brain slice.²²¹ Following the selective excitation that generates a gradually decaying MR signal, the gradient along the z-axis is switched off. During the signal sampling, two additional gradients are then switched on sequentially: one along the x-axis (the so-called frequency-encoding gradient) and the other along the y-axis (the so-called phase-encoding gradient). These two additional gradients encode the spatial locations across the slice, thus parcelling it into voxels. Depending on how the timing, strengths, directions, and durations of the frequency-encoding and the phase-encoding gradients are combined, MR signals with different characteristics can be sampled.²²²

However, the major caveat is that the scanner cannot sample the signals directly from individual voxels but can only measure the cumulative signal from the entire slice.²²³ Hence, to disentangle the contributing signals from individual voxels and reconstruct their relative spatial locations, the cumulative signal has to be mathematically broken down into its constitutive components. Referred to as the image reconstruction, the mathematical decomposition of the cumulative signals in contemporary MR scanners is performed by an automated computer algorithm called the *inverse fast Fourier transform*.²²⁴ The consequence for the imaging is that a single spatially encoded signal measurement does not suffice for the reconstruction of one 2D slice. Instead, it is necessary to collect many signal measurements, referred to as data points. Furthermore, these data points have to be encoded through a particular temporal sequence of different combinations of x- and y-gradients so that the *inverse Fourier transform* can use the resulting set of signals to create a single 2D image.²²⁵ Such

220 Huettel, Song, and McCarthy, *Imaging*, 91–93. Historically, the method of slice selection was initially described in 1974. See Garoway, Grannell, and Mansfield, “Selective Irradiative Process.”

221 The resulting cross-sectional MRI image is referred to as a 2D slice since its ‘thickness’ is one voxel, which is the smallest spatial unit. Yet, depending on the spatial resolution of the 2D slice, a single voxel stands for a brain slice whose physical thickness can range from less than a millimetre to several millimetres.

222 For details on slice selection, as well as phase and frequency encoding, see Huettel, Song, and McCarthy, *Imaging*, 91–97.

223 Huettel, Song, and McCarthy, 93.

224 Haacke et al., *Imaging*, 240. The basic principle underlying the *fast Fourier transform* is the powerful analytical method developed by the nineteenth-century French physicist and mathematician Joseph Fourier. According to this method, any complex signal can be described as a weighted sum of simple waves of various wavelengths and amplitudes. For a succinct description, see Jezzard and Clare, “Principles,” 78–80. Whereas the *forward Fourier transform* “can convert image-space data into k-space data,” the *inverse Fourier transform* can convert k-space data into an image. Huettel, Song, and McCarthy, *Imaging*, 110. As a mathematical tool, the *Fourier transform* was originally developed for continuous analogue values. However, in contemporary scanners, the raw MR signals are converted from continuous analogue values into digital discretised data immediately upon sampling. Haacke et al., *Imaging*, 299. The *discrete Fourier transform* is a specifically modified version of *Fourier transform*, which is applicable to digital data. The *fast Fourier transform*, in turn, is “an efficient computer algorithm for calculating the *discrete Fourier transform*.” Haacke et al., 240.

225 Huettel, Song, and McCarthy, *Imaging*, 110. For the initial articles that introduced the *Fourier transform* approach to MRI imaging in the mid-to-late 1970s, see Kumar, Welti, and Ernst, “NMR

a targeted collection of data points is achieved by implementing a specific notation scheme called raw-data space or k-space.

Simply formulated, k-space is a way of collecting, organising, and storing the MR signal measurements (i.e., raw data) so that the standard mathematical reconstruction algorithm can translate them, without any information loss, into a 2D image (fig. 3.4).²²⁶ A particular coordinated step-by-step acquisition of MR signals necessary to produce a single slice with the desired characteristics is referred to as filling k-space.²²⁷ Crucially, the “degree to which the image [of the brain] is faithfully reconstructed depends on the completeness of k-space coverage.”²²⁸ Strictly speaking, each point in k-space corresponds to a single measurement of a total MR signal from an entire cross-section of the brain at a specific point in time and under the implementation of a particular combination of magnetic gradients.²²⁹ Thus, as a collection of raw MR data, k-space is built up of individual signal measurements that are “suitably stacked” to allow their subsequent mathematical transformation into a cross-sectional 2D image.²³⁰ Since each data point is designated by a numerical value, k-space is a mathematical entity called a matrix—an array of numbers arranged in a grid of rows and columns.²³¹

After the k-space matrix has been filled, the *inverse fast Fourier transform* translates it into a new matrix—a 2D MR image. In the resulting image matrix, the individual numbers designate the relative strengths of the MR signals across different voxels. The rows and columns of this matrix are assigned those sets of spatial coordinates that the scanner had mapped onto the brain during the measurement. Therefore, expressed in mathematical terms, a 2D MR image describes how the relative signal intensity changes as a function of x and y spatial coordinates.²³² By assigning a particular grey-scale value to each number in the new matrix, the algorithm automatically visualises this image as a grey-scale brain slice (see fig. 3.3).

Based on my analysis, two aspects should by now be apparent. First, even though k-space and the image reconstructed from it contain the same informational content, “there is *not* a one-to-one relationship between points in k-space and voxels in image

Fourier Zeugmatography”; Mansfield and Maudsley, “Line Scan”; and Mansfield and Maudsley, “Medical Imaging by NMR.”

226 In other words, k-space representation and the image reconstructed from it have the identical informational content. The difference between k-space and the image is the physical units used to express this informational content. The information within k-space is non-spatially encoded—the unit is not distance but spatial frequency (1/distance). See Mansfield, “Snap-Shot MRI,” 269.

227 Huettel, Song, and McCarthy, *Imaging*, 90.

228 Haacke et al., *Imaging*, 308.

229 More precisely, a filled k-space offers a mathematical description of how the measured MR signal changed depending on the magnitude of magnetic field gradients over time. For more details on k-space, see Huettel, Song, and McCarthy, *Imaging*, 109–17.


230 Mansfield, “Snap-Shot MRI,” 269.

231 In the visualisation of k-space seen in fig. 3.4, the computer software had already automatically attributed various shades of grey to each numerical value comprising the k-space matrix.

232 The relative signal intensity, in turn, is determined by the relative density of protons characterised by the physical property that has been highlighted through the targeted sampling. See Mansfield, “Snap-Shot MRI,” 269.

space.”²³³ Instead, each point in k-space, as a signal measured from the entire slice, contributes to each voxel of the reconstructed 2D slice. And second, considering the numerous individual signal measurements and subsequent transformations that underpin its production, it can be argued that even a single 2D grey-scale brain slice is already a composite image.

Figure 3.4. 2D fMRI slice visualised in k-space (i.e., raw-data space) format.




K-space, however, is not only a concrete way of storing and organising the data points acquired. At a more abstract level, k-space is also a mathematical framework that informs the entire process of data acquisition. In this latter sense, k-space governs the temporal organisation of all individual steps that comprise spatial encoding and signal sampling to allow an optimal translation of the brain's properties of interest into MR images with desired characteristics.²³⁴ Hence, k-space can be sampled in many different ways. This depends on how many data points are collected and which chronological sequence of RF pulses and time-based gradients of a particular strength,

233 Huettel, Song, and McCarthy, *Imaging*, 114 (emphasis in original). It is also interesting to note how different parts of k-space, which encode different spatial frequency components of the image, contribute to the resulting image space. The centre of k-space “provides low-spatial-frequency information, retaining most of the signal but not fine details.” Ibid. This is because the point at the centre of k-space was measured when the signal contributions from all voxels were at the same phase, resulting in the maximum total signal. Peripheral parts of the k-space provide “high-spatial-frequency information, and thus more image detail” without contributing much signal to the image. Ibid.

234 See Haacke et al., *Imaging*, 139–330.

duration, and spatial orientation was used to encode the signals.²³⁵ A specific temporal combination of time-dependent gradients and RF pulses is called a k-space trajectory and can be represented in the form of a diagram (fig. 3.5).

Figure 3.5. Diagrammatic visualisations of different k-space trajectories. Left: conventional trajectory used in the acquisition of a structural 2D slice. Right: single continuous trajectory that underlies the acquisition of a 2D fMRI slice with the EPI sequence. From: Stehling, Turner, and Mansfield, “Echo-Planar Imaging,” 44, fig. 1, C and D. ©American Association for the Advancement of Science.



Crucially, the choice of a particular k-space trajectory does not only determine the speed of signal acquisition. It also informs all spatial features of the resulting image, such as its size and resolution, as well as the presence and distribution of various distortions.²³⁶ Thus, diverse k-space trajectories can be applied flexibly and selectively to translate one and the same brain into image spaces with entirely different characteristics. This is precisely what happens during the production of high-resolution structural and low-resolution functional images. The spatial features of these two imaging modalities—defined by the different sizes and numbers of voxels constituting them—are constructed through distinctly different k-space trajectories.

For example, to generate a structural 2D slice, k-space is filled line by line, following the application of a separate RF pulse for each line (fig. 3.5, left).²³⁷ Consequently, producing a single structural slice requires several hundred separate RF pulses, and it takes a long time to fill its k-space. By contrast, in the high-speed imaging sequence called EPI (echo-planar imaging), which is the standard method of acquiring fMRI images, signal generation and data sampling are organised differently.²³⁸ Specifically, the EPI sequence is based on an “unconventional” continuous

235 Via the aforementioned quantum mechanical signal equation, k-space offers a mathematical description of how the MR signal changes as a function of the magnitude of the phase- and frequency-encoding gradients over time. See Mansfield, “Snap-Shot MRI,” 269.

236 Huettel, Song, and McCarthy, *Imaging*, 113–20.

237 For details regarding the basic k-space trajectory in anatomical imaging, see Huettel, Song, and McCarthy, 112–13.

238 See Huettel, Song, and McCarthy, 147–54.

k-space trajectory “in which alternating lines are scanned in opposite directions.”²³⁹ This highly efficient back-and-forth trajectory, which necessitates fast switching of the phase- and frequency-encoding magnetic field gradients, enables the sampling of hundreds of data points following a single RF excitation (fig. 3.5, right).²⁴⁰ As a result of this specific k-space trajectory, all data points required to reconstruct an entire 2D fMRI slice can be acquired within a fraction of a second.

The speed and accuracy with which image acquisition has to be performed in contemporary clinical and research settings have led to the computerisation, standardisation, and automation of the processes described above. Thus, nowadays, these processes are black-boxed behind algorithm-based pulse sequences. Researchers use commands of the graphic user interface to set in motion computer algorithms, which then automatically govern the generation, spatial encoding, and sampling of MR signals, as well as the subsequent image reconstruction. Researchers can only indirectly control these cascades of transformations by choosing among various predefined parameter options listed in the user interface on their computer screen. This nevertheless allows them to shape the measurement with a considerable level of flexibility by making decisions about an entire spectrum of parameters. These parameters include the type of the pulse sequence, size of the image matrix, number of slices, and their thickness, as well as the size of gaps between consecutive slices.

A 3D fMRI image usually comprises twenty to fifty slices, each with a matrix of 64 x 64 or 128 x 128 voxels, whose size ranges from 3 x 3 x 3 mm to 4 x 4 x 4 mm.²⁴¹ Conversely, a structural 3D image comprises on average almost two hundred slices with an underlying matrix of either 256 x 256 or 512 x 512 voxels, whose size is often as small as 1 x 1 x 1 mm.²⁴² Furthermore, although slices are by convention acquired perpendicular to the z-axis, researchers can perform the imaging in any desired spatial orientation. Since these decisions have a decisive impact on the properties of the images thus generated, all the parameters employed during the acquisition have to be mentioned explicitly in the published fMRI study.

On the whole, my analysis in this and the previous section has made evident the multiple steps of the precisely coordinated physical interventions and advanced mathematical

239 Huettel, Song, and McCarthy, 148. This particular k-space trajectory was developed by the English physicist Peter Mansfield in 1977. For a succinct description, see Mansfield, “Snap-Shot MRI,” 266–70. For the original article, see Mansfield, “Spin Echoes.” Interestingly, the implementation of Mansfield’s elegant mathematical concept was technically demanding, as it requires strong magnetic field gradients that have to be switched on and off very rapidly. Hence, not before the early 1990s did MR scanners that could implement an EPI sequence enter the medical market. See Stehling, Turner, and Mansfield, “Echo-Planar Imaging,” 48. For his development of the EPI sequence, Mansfield shared the 2003 Nobel Prize in Physiology or Medicine with Paul Lauterbur.

240 For details, see Stehling, Turner, and Mansfield, “Echo-Planar Imaging,” 44–45.

241 See Baek et al., “Impaired Awareness,” 1626; and de Lange, Roloefs, and Toni, “Self-Monitoring,” 2053.

242 See Baek et al., “Impaired Awareness,” 1626; and de Lange, Roloefs, and Toni, “Self-Monitoring,” 2053.

modelling necessary to produce functional and structural imaging data. Through these steps, the features of interests of the active brain are selectively translated into a hybrid object, which is, at the same time, an image and a mathematical entity. Thus, I have shown that the mathematical and visual aspects of functional imaging data are not antithetical—as Anne Beaulieu has claimed—but instead mutually entangled.²⁴³ We have seen that the essential visual aspects of an MRI image—such as its spatiality, contrast, and resolution—result from pervasive mathematical structuring that underlies all aspects of the measurement. Specifically, I have demonstrated how the very ‘imageness’ of the structural and functional imaging data is constructed through a tailored use of the quantum mechanical signal equation, the Cartesian coordinate system, matrix algebra, and the *inverse Fourier transform*.

All these diverse mathematical models are inscribed into the data and enmeshed with the visualised properties of the measured brain. Yet, this substantial mathematical modelling does not mean that the visual aspects of the fMRI data are subordinated to the numerical. In fact, as we will see later in this chapter, the ability of the imaging data to be transformed back and forth from the numerical into various visual forms plays a pivotal role during their subsequent processing. But before moving on to discuss various stages of fMRI data processing, there is one additional aspect of the data acquisition that we first need to address—its temporal dimension.

3.2.3 Acquiring 4D BOLD Imaging Datasets

Brain activity and the correlated physiological changes unfold not only in space but also in time. In contrast, brain anatomy remains constant throughout the fMRI image acquisition. For this reason, the temporal dimension of the measurement has different imports on structural and functional images. This, in turn, has consequences on how these two types of imaging data are acquired and subsequently analysed. Moreover, as I will show in what follows, it also has consequences on why, when visualised upon the finished acquisition, structural data are immediately visually legible,²⁴⁴ whereas functional data are not.

Since structural images provide information about static anatomical features, it suffices for researchers to collect a single 3D volume—a set of up to two hundred 2D slices that cover the entire brain. As discussed previously, this single volume typically has a high spatial resolution as it is built up of several million very small voxels.²⁴⁵

243 Based on a series of interviews, Anne Beaulieu has concluded that neuroscientists tend to downplay the visual aspects of functional neuroimaging data while emphasising their quantitative character. Beaulieu has criticised this tendency, calling it iconoclastic. Yet, at the same time, she has claimed that functional scans are characterised by an irreconcilable tension between their visual and numerical aspects since the measurements have to be displayed in space. According to Beaulieu, the construction of the spatiality of functional brain scans necessitates the implementation of pictorial conventions. See Beaulieu, “Not the (Only) Truth,” 53–86.

244 As defined earlier, in my use, a legible image is one in which the information of interest is codified in such a way that, at least in principle, it can be accessed by visual inspection of that image.

245 For instance, a structural volume comprised of 176 slices with an underlying matrix of 256 x 256 contains more than eleven million voxels.

Collecting sufficient raw data to reconstruct such a spatially detailed 3D image requires a long acquisition time of up to ten minutes.²⁴⁶ Conversely, since fMRI data measure rapidly and continually changing neurophysiological processes, they have to be collected quickly and repeatedly throughout the experiment. The required speed, however, can only be achieved at the expense of spatial detail. Therefore, fMRI slices are built of relatively large voxels—ca. 3 x 3 x 3 mm—each containing several million neurons treated as a single spatial unit.²⁴⁷ Due to its coarser spatial resolution, an entire fMRI volume consisting of twenty to forty slices can be acquired every 1–3 seconds over a period of twenty to thirty minutes.²⁴⁸ Consequently, functional acquisition produces what is called a time series—a collection of 3D fMRI images that have been generated at regularly spaced periods. An fMRI dataset is thus characterised by one temporal and three spatial dimensions.

Upon finished measurement, both structural and functional data are automatically visualised on the computer screen in spatial form as two separate series of grey-scale cross-sectional brain slices. Although, as discussed previously, BOLD fMRI slices primarily encode the relative concentration of deoxygenated haemoglobin across the brain, they also contain some additional but very rudimentary anatomical information.²⁴⁹ Yet, due to the particular acquisition method through which fMRI slices were generated, the additional anatomical information is not clearly encoded in these images. Hence, even experts who possess the visual skills necessary to ‘read’ structural brain images by knowing how to distinguish distinct anatomical features cannot “identify boundaries between different types of tissue” in BOLD fMRI slices.²⁵⁰

This means that in terms of anatomy, BOLD fMRI slices are illegible or, in other words, not clear enough to be read.²⁵¹ However, neuroimaging operates under the premise that different anatomical structures have different specialised functions.²⁵² For this reason, it is crucial to establish not just the spatial location of the task-induced brain activity (i.e., its Cartesian coordinates) but also its anatomical location. Consequently, in every fMRI study, prior to the stage of functional imaging, a high-resolution structural brain volume has to be acquired, with the subject keeping the same position within the scanner for both types of imaging sequences. As we will see later in this chapter, it is by using the resulting structural dataset that researchers construct the anatomical legibility of fMRI images.

246 Ashby, *Statistical Analysis*, 4.

247 Logothetis, “What We Can Do,” 875. For instance, a BOLD fMRI volume comprised of 32 slices with an underlying matrix of 64 x 64 contains about 131,000 voxels.

248 Logothetis, 875.

249 Jenkinson and Chappell, *Neuroimaging Analysis*, 170.

250 Huettel, Song, and McCarthy, *Imaging*, 248. I am using the term ‘reading’ here in the sense introduced by Sybille Krämer. It denotes the learned ability to overlook the epistemically insignificant visual features while also knowing which of the few relevant visual features to focus on to obtain the information of interest encoded in the image. See Krämer, “Operative Bildlichkeit,” 102.

251 As stated in the introduction to this chapter, in my designation, illegibility designates an intrinsic property of some images.

252 See section 2.3.1.

But an even bigger challenge arises from the fact that fMRI data are also illegible in a more categorical way. What I mean is that the primary information of interest—i.e., determining in which voxels the task-induced brain activity occurred—remains inaccessible to visual inspection of BOLD slices. Put differently, whereas an expert can look at structural MRI slices and clearly identify different anatomical structures in them, the same expert cannot disambiguate between active and inactive voxels by visually examining fMRI data.²⁵³ There are several distinct reasons why even experts cannot ‘read’ functional datasets. As I am about to show, all of these reasons are a direct consequence of the inherent limitations of the fMRI technology.

When viewed as a collection of slices, fMRI data are illegible because this mode of visualisation foregrounds their spatial features at the expense of their temporal dimension. Yet, the temporal dimension is a crucial aspect of an fMRI dataset. This is because BOLD fMRI does not use a static amount of deoxygenated haemoglobin as the indicator of neural activity in a particular voxel, but instead a distinct temporal change in the concentration of these molecules. This temporal change, in turn, arises from what, in specialist terms, is known as neurovascular coupling. Neurovascular coupling is a phenomenon entailing the interplay of multiple physiological reactions that accompany neural activity, such as changes in the local blood flow, blood volume, and oxygen consumption.²⁵⁴ Notably, the exact nature and the underlying physiological mechanism of this dynamic phenomenon remain poorly understood.²⁵⁵ It has been experimentally shown that with a several seconds delay in relation to the neural activity, which lasted only a few milliseconds, a temporary oversupply of oxygenated blood at the capillaries around the active neurons takes place.²⁵⁶ Since the changes in the blood flow are of much higher intensity than the local oxygen consumption, the relative amount of oxygenated haemoglobin at first increases. But then, following further changes in the local blood flow and metabolism, the relative amount of oxygenated haemoglobin gradually declines. After a while, it returns to the baseline level it had before the onset of the neural activity at the given location.

As a result of these interrelated metabolic and vascular processes, the MR signal from an active voxel begins to rise with a delay of 1–2 seconds, reaches a peak at about 4–6 seconds and returns to baseline by 12–20 seconds after the correlated neural

253 Strictly speaking, it is not the voxel as an elementary spatial unit of the images that is active or inactive, but a part of the brain to which it refers. In this metonymic sense, the expressions ‘active voxel’ and ‘inactive voxel’ are regularly used in neuroscientific literature. See, e.g., Huettel, Song, and McCarthy, *Imaging*, 357. It is in this sense that I use these terms throughout this chapter. I have chosen to adopt these terms because they pertinently draw attention to the fact that, in an fMRI study, any claim researchers make about the task-induced brain activity is necessarily mediated through the imaging data. As we will see by the end of this chapter, although researchers make inferences about active brains, they do so by searching for traces of the neural activity of interest in the imaging data.


254 See Bandettini et al., “Time Course,” 390–97; and Kwong et al., “Human Brain Activity,” 5675–79.

255 For a detailed account, see Buxton, *Functional Magnetic Resonance Imaging*, 5–63. See also Logothetis, “Neural Basis,” 1008–30; and Logothetis, “What We Can Do,” 869–78.

256 For an overview, see Huettel, Song, and McCarthy, *Imaging*, 196–200.

activity took place.²⁵⁷ Importantly, it has been experimentally determined that the height of the peak and the relative timing of the signal's rising and falling varies across individual subjects and even from one anatomical brain area to another.²⁵⁸ Despite such variations, the basic underlying pattern nevertheless remains the same. This pattern can be mathematically modelled and visualised in the form of a curve with a specific shape and is called the haemodynamic or BOLD response (fig. 3.6).²⁵⁹ Crucially, this distinct pattern in which the MR signal from a single voxel changes over time due to neurovascular coupling is the information of interest in fMRI images.²⁶⁰ Thus, BOLD fMRI uses a delayed phenomenon that extends over many seconds as a proxy for the neural activity that lasts only a fraction of a second.

Figure 3.6. Canonical basis function that mathematically models the haemodynamic response. From: Ashburner et al., "SPM12 Manual," 245, fig. 31.9. ©Wellcome Centre for Human Neuroimaging, London.



However, since a single fMRI slice encodes the values of MR signals at specified locations at a given instant, it can only contain a temporal fragment of the BOLD

257 Poldrack, Mumford, and Nichols, *Handbook*, 70–72.

258 Aguirre, Zarahn, and D'Esposito, "Variability of Hemodynamic Responses," 360–69.

259 See Poldrack, Mumford, and Nichols, *Handbook*, 72.

260 Huettel, Song, and McCarthy, *Imaging*, 208–14.

response. This means that the information of interest—whether a chosen task elicited a BOLD response or not—is not encoded within a single slice or even a single volume. Instead, the BOLD response is spread across a sequence of 3D images acquired at different time points. For this reason, the measurement has to be performed in a series of multiple ultra-quick scans that repeatedly sample the entire brain. The sampling rate determines the data's temporal resolution and is one of the parameters that researchers can specify. Due to the relatively long duration of the BOLD response, it suffices to sample a single brain volume every 1–3 seconds.²⁶¹

Based on my analysis so far, it might appear as if fMRI data are visually illegible due to their spatial visualisation that effectively sidelines their temporal dimension. Were this the case, the problems could be solved easily. A 4D dataset can also be visualised in a way that foregrounds its temporal aspects, albeit thus necessarily omitting the spatial relationships among the individual voxels. In its temporal form, a 4D dataset is displayed as a series of curves, one for each voxel.²⁶² Such a curve, called a time course, shows how the MR signal at a chosen voxel changes over time (fig. 3.7). Since we know by now that the BOLD response has a distinct visual shape, we might assume that researchers could pinpoint the presence of brain activity of interest by visually inspecting such time courses. Yet, this is not the case for two distinct reasons.

First, even when visualised as a set of time courses, BOLD data remain illegible to visual inspection because they contain a massive amount of noise. In other words, the haemodynamic response that causes a temporary decrease in the local concentration of deoxygenated blood is only one of many factors that cause a measurable change in the MR signal.²⁶³ Normal physiological processes such as breathing and heartbeat, the ongoing brain activity of no interest, the subject's minimal head movements, and a variety of potential technical problems with the scanner are some of the many additional factors that introduce noise into the data. As a matter of fact, the contributions that are considered non-meaningful for an fMRI study make up 90–99% of the signal changes measured.²⁶⁴

Moreover, based on my earlier analysis,²⁶⁵ it is easy to conclude that both the task-induced BOLD responses and noise are indexically inscribed into the imaging data through the measurement. As such, BOLD responses and noise are indiscernible from one another, except through statistical analysis. Hence, to counter the problem of noise in the data, all experimental task conditions are presented not once but instead repeated many times during the data acquisition. This repetition allows researchers to later statistically average non-meaningful changes in MR signals across multiple trials and thus filter them out.²⁶⁶ Consequently, an fMRI experiment must generate many hundreds of 3D fMRI images to enable the production of even a single functional

261 For a detailed account, see Huettel, Song, and McCarthy, 220–29.

262 Huettel, Song, and McCarthy, 211.

263 For an overview of various sources of noise in fMRI, see Huettel, Song, and McCarthy, 255–67.


264 Amaro and Barker, "Study Design," 222. For the relation between signal and noise at different field strengths, see Krüger, Kastrup, and Glover, "1.5 T and 3.0 T," 595–604.

265 See section 3.2.1.

266 See Dale and Buckner, "Selective Averaging," 329–30.

brain map. For example, in the de Lange, Roelofs, and Toni study, for each patient, the researchers acquired 547 brain volumes, each of which consisted of 32 individual functional slices.²⁶⁷ They then repeated this procedure for each of their eight study participants.

Figure 3.7. BOLD MRI signal from a single voxel visualised in the form of a time course. The dots designate the individual sampling points.



The second reason for the illegibility of time courses is that fMRI can measure neither the task-induced neural activity nor its correlated BOLD response in absolute but only in relative terms.²⁶⁸ This is because, as discussed earlier, to draw conclusions about the task-induced activations, researchers must calculate the differential BOLD responses across two or more experimental conditions. But, since these conditions mutually alternate throughout the measurement, the differential responses are spread across different time points. Hence, the task-induced differential responses cannot be spotted through a visual inspection of time courses that display a linear succession of individual time points at a single voxel (see fig. 3.7).

To conclude, I have underscored that BOLD fMRI slices are a product of many mutually nested transformations that indexically link these images to the experimental subjects' active brains. Even so, based on my analysis, I argue that BOLD slices cannot be understood as indexical signs of neural activity because they are essentially illegible and, therefore, in the context of an fMRI study, still meaningless. In the remainder of this chapter, I will show that the semantic potential of BOLD fMRI data has to be articulated through statistical analysis, which makes the informational content of these images accessible by isolating it from noise and synthesising it across temporally spaced experimental conditions and spatially distinct voxels. We will see that during this process, illegible BOLD fMRI data are gradually transformed into visually legible functional brain maps.

267 See de Lange, Roelofs, and Toni, "Self-Monitoring," 2053.

268 Huettel, Song, and McCarthy, *Imaging*, 354.

Hence, on the whole, BOLD fMRI images are best understood as intermediary inscriptions whose function is, first and foremost, to bridge the otherwise insurmountable gap between the subject's active brain and the functional maps. As the output of the measurement procedure, BOLD fMRI images have a fixed material form. Owing to this fixed material form, they can be archived, copied and transported, shared within the scientific community and even reused in later studies.²⁶⁹ However, as the following sections will make evident, the key feature of fMRI images is their mutability, which arises from the fact that various mathematical operations can be performed on them. Owing to their mutability, these images are able to fulfil their primary epistemic function as the working material for subsequent transformations. In what follows, we will examine these transformations and discuss their epistemic implications.

3.3 Preprocessing: Constituting the Analysability of fMRI Data

Having collected the imaging data for all their study participants, researchers then move on to the subsequent stages of the experiment, during which they process the raw datasets. Across these stages, researchers aim to translate the illegible and noisy fMRI datasets into visually accessible functional brain maps. Called the processing pipeline, this procedure entails a sequence of algorithmic steps that systematically address various types of noise. In the following sections, I will examine these steps by focusing on how researchers make judgments about what counts as noise in their data and which operations they perform to remove it. I will show that by making these judgments, researchers inscribe a range of both explicit and implicit theoretical assumptions into the imaging data. It is important to unpack these assumptions since they are invisible in the functional maps as the products of the analytical pipeline. Yet, although invisible, these assumptions inform the maps' potential scientific validity and their ability to produce new insights into hysteria or, at a more general level, any other phenomenon under study.²⁷⁰

Generally speaking, a processing pipeline comprises two distinct stages. Each stage is tailored to deal with a specific type of noise—random or systematic. The primary sources of random noise in an fMRI experiment include, first, brain processes unrelated to the experimental task, and second, variations in how the subjects performed the task at hand.²⁷¹ This type of noise is study-specific because it depends on the concrete experimental task and the subjects selected. To remove it, researchers deploy statistical analysis during the main stage of processing. But before statistical analysis can be

269 As discussed previously, the underlying structure of each slice is a matrix—an array of numbers arranged in rows and columns.

270 To demonstrate the analytical variability of fMRI processing pipelines, one meta-study focused on ten standard preprocessing and modelling steps. By considering between two and four default options for each step and then taking into account their various combinations, the authors arrived at 6,912 different pipelines. When applied to the same dataset, each pipeline resulted in a different functional map. See Carp, "Analytic Flexibility."

271 How a task is performed varies not just among different subjects but also over a single subject's repeated trials during the experiment. See Huettel, Song, and McCarthy, *Imaging*, 262.